arXiv:1405.6899v1 [cs.MA] 27 May 2014

Big, but not unruly: Tractable norms for anonymous game structures

Truls Pedersen[1], Sjur Dyrkolbotn[2], and Piotr Kaźmierczak[1,3⋆]

[1] Dept. of Information Science and Media Studies, University of Bergen, Norway
[2] Durham School of Law, Durham University, UK
[3] Dept. of Computing, Mathematics and Physics, Bergen University College, Norway
truls.pedersen@infomedia.uib.no, s.k.dyrkolbotn@durham.ac.uk, phk@hib.no

Abstract. We present a new strategic logic NCHATL that allows for reasoning about norm compliance on concurrent game structures that satisfy anonymity. We represent such game structures compactly, avoiding models that have exponential size in the number of agents. Then we show that model checking can be done in polynomial time with respect to this compact representation, even for normative systems that are not anonymous. That is, as long as the underlying game structures are anonymous, model checking normative formulas is tractable even if norms can prescribe different sets of forbidden actions to different agents.

1 Introduction

Logics of strategic ability such as Alternating-time Temporal Logic (ATL) [4] or Coalition Logic [13] have gained much interest in the multi-agent systems community in recent years. The language of ATL (of which Coalition Logic is the next-time fragment) allows for expressing formulas about strategic ability of (coalitions of) agents, and it has been used for modelling open multi-agent systems [3]. Originally, ATL was used for modelling *heterogeneous* systems. Recently, however, a semantics for ATL tailored towards systems exhibiting some degree of *homogeneity* was presented [15,14].

In this paper, we continue this line of research, noting that the *homogeneity* requirement relied on in [15] has been studied independently in game theory, where *anonymity* is the name given to a corresponding property of a normal form game, which is obtained when payoff functions remain invariant under permutations of players, see e.g., [5,12,6].

We tackle the question of regaining some of the expressive power lost by requiring anonymity, and we do so by using *normative systems*. These have emerged as a promising and powerful framework for coordinating multi-agent systems [16,17,10,1,8,2]. They allow the modeller to constrain the behaviour of agents, and can thus provide a way to ensure that the global behaviour of the system exhibits some desirable properties. We point out that normative systems

⋆ Piotr Kaźmierczak's research was supported by the Research Council of Norway project 194521 (FORMGRID).

the way we understand them are sometimes also called *social laws*, and are simply behavioural restrictions on agents developed by an offline designer (who is not part of the model), much in the spirit of Shoham & Tenneholtz's seminal paper [17], and thus different from normative systems known from deontic logic literature, since we abstract away from things like obligations, institutions, etc.

A key issue is the question of *compliance*. Even if a normative system is *effective* in the sense that it will ensure that the objective holds, under the assumption that all agents comply with it, interesting questions and increased expressive power arise when one assumes that only *some* agents comply. Despite the anonymous settings it is not irrelevant *who* those agents are, in particular, and as a consequence, a normative system provides us with a way to regain expressive power that is lost by imposing anonymity. We also show that doing so for anonymous game structures is possible while maintaining the compact representation and the tractable model checking that comes with it.

In short, our contribution in this paper combines the following four different themes: strategic logic ATL, normative systems, homogeneous structures and anonymous games. The resulting *Norm Compliance Homogeneous Alternating-time Temporal Logic* (NCHATL), in particular, arises from adding norms to *homogeneous* ATL [14] in a way that renders the resulting model checking problem polynomial in the number of agents.

The structure of the paper is as follows. In Section 2 we introduce the formal background, recalling the definition of concurrent game structures (CGSs) and the definition of anonymity used in game theory and social choice theory. We also present a special case of the construction used in [15], showing how an anonymous CGS can be succinctly represented as a concurrent game structure with roles (an RCGS) where the number of roles is exactly one. We go on to formulate the notions of *norms* and *norm compliance* as they are used in the multi-agent systems community. Then in Section 3 we define a semantics for NCHATL and investigate model checking for this logic, showing that it is tractable. We conclude in Section 4.

2 Formal Background

We start by introducing some definitions of the formal framework used in the paper. The logical language we use, $\mathcal{L}_{\text{NCHATL}}$, is based on ATL [4], extended with one extra operator that we use to express norm compliance. Formally, the language is generated by the following BNF:

$$\varphi ::= \top \mid p \mid \neg\varphi \mid \varphi \vee \varphi \mid \langle\langle C\rangle\rangle \bigcirc \varphi \mid \langle\langle C\rangle\rangle \Box\varphi \mid \langle\langle C\rangle\rangle \varphi \mathcal{U} \varphi \mid \langle C\rangle\varphi$$

where p is a propositional symbol, and C is a coalition of agents.

The language of NCHATL contains three types of modalities:

- $\bigcirc$, $\Box$ and $\mathcal{U}$ are standard *temporal* operators known from many temporal logics, and stand for "next state", "some future state" and "until", respectively;

– $\langle\langle C \rangle\rangle$ is a *strategic ability* operator, and its intuitive meaning is that the coalition $\langle\langle C \rangle\rangle \bigcirc \varphi$ has a *joint strategy* for enforcing a formula φ in the next state;
– finally $\langle C \rangle$ is the *norm compliance* operator, which intuitive reading is that the coalition C has a strategy to achieve φ if all its members comply to a given normative system.

2.1 Anonymity

We now define *concurrent game structures* known from [4] and used for ATL interpretation, and formalize the anonymity requirement mentioned in the introduction. We then define the compact representation of such structures, which provides the backbone for the semantics of NCHATL.

Definition 1 (Concurrent Game Structure). *A* CGS *is a tuple* $S = \langle \mathcal{A}, Q, \Pi, \pi, \mathbb{A}, \delta \rangle$ *where:*

– $\mathcal{A}$ *is a non-empty set of players. In this text we assume* $\mathcal{A} = [n]$ *for some* $n \in \mathbb{N}$*, and we reserve* n *to mean the number of agents.*[4]
– Q *is the non-empty set of states.*
– Π *is a set of propositional letters and* $\pi : Q \to \wp(\Pi)$ *maps each state to the set of propositions true in it.*
– $\mathbb{A} : Q \times \mathcal{A} \to \mathbb{N}^+$ *is the number of available actions in a given state. We also say that for each state* $q \in Q$ *a* move vector *is a tuple* $\langle \alpha_1, \ldots, \alpha_k \rangle$ *s.t.* $1 \leq \alpha_a \leq \mathbb{A}_a(q)$ *for each* $a \in \mathcal{A}$*. D is then a* move function *which given a state* $q \in Q$ *outputs a set of move vectors.*
– *For each* $q \in Q$ *and a move vector* $\langle \alpha_1, \ldots, \alpha_k \rangle \in D(q)$ *a* transition function *produces a state* $\delta(q, \alpha_1, \ldots, \alpha_k) \in Q$ *which is a successor of* q *when every agent* $a \in \{1, \ldots, k\}$ *chooses* α_a*.*

Inspired by the corresponding notion from game theory [7,5,6], we will say that a CGS S is *anonymous* if and only if:

$$\forall_{q \in Q, i, j \in \mathcal{A}}, \mathbb{A}(q, i) = \mathbb{A}(q, j) \text{ and}$$
$$\forall_{q \in Q, i, k \in \mathcal{A}}, \delta(q, \ldots, \alpha_i, \ldots, \alpha_k, \ldots) = \delta(q, \ldots, \alpha_k, \ldots, \alpha_i, \ldots)$$

Any anonymous CGS can be represented compactly as an RCGS – a Concurrent Game Structure with Roles [15]. In fact, the class of anonymous CGS's corresponds to the class of RCGS's with a single role, a simplified definition of which can be given as follows [14].

Definition 2. *A* 1RCGS *is a tuple* $R = \langle \mathcal{A}, Q, \Pi, \pi, \mathbb{A}, \delta \rangle$ *where*

– $\mathcal{A}$*,* Q*,* Π*, and* π *are defined as in Definition 1,*
– $\mathbb{A} : Q \to \mathbb{N}^+$ *is the number of available actions in a given state.*

[4] For the sake of brevity, we use the notation $[n]$ to indicate the set of numbers $1 \leq i \leq n$.

– *For every state we have a set of vectors* $P(q) = \{F \in [n]^{[\mathbb{A}_q]} \mid \sum_{i \leq \mathbb{A}_q} F_i = n\}$. *We will refer to the elements of* $P(q)$ *as the* profiles *at* q. *For every state* q *and every such profile* $F \in P(q)$ *we have a* successor state $\delta(q, F) = q'$.

The profiles assign a natural number to each action such that the sum of these numbers (over all actions) sums up to the number of agents n. The intended meaning is that the profile describes how many agents perform each action. We also define partial profiles at $q \in Q$, for all $A \subseteq \mathcal{A}$ as follows:

$$P(q, A) = \left\{ F \in [n]^{[\mathbb{A}_q]} \;\middle|\; \sum_{i \leq \mathbb{A}_q} F_i = |A| \right\}.$$

It is not hard to see that a 1RCGS can be given to provide a succinct representation for any CGS which satisfies the anonymity requirement; as the permutations of the action profiles are irrelevant, we only need to record *how many* agents performed each action.

2.2 Normative systems

Following [2,10] we define a normative system as a map $\eta : Q \times \mathcal{A} \to 2^{\mathbb{N}^+}$, giving, for each state and agent, the set of actions that are forbidden for that agent in that state. We require $\eta(q, a) \in [\mathbb{A}_q]$ for all $q \in Q, a \in \mathcal{A}$ and that η is such that, for every state, there is at least *some* legal action. That is, $\forall q \in Q, a \in \mathcal{A} : [\mathbb{A}_q(a)] \setminus \eta(q, a) \neq \emptyset$.

To account for "disobedience" of certain agents (i.e. those that do *not* comply with a normative system), we consider normative systems *restricted* to specific coalitions. Such a restriction means that only actions that are controlled by a given coalition are blacklisted (intuitively, anyone not belonging to that coalition is free not to comply with the normative system). We use the $\restriction$ symbol to denote such restrictions, and formally define it below:

$$(\eta \restriction C)(q, a) = \begin{cases} \eta(q, a) & \text{if } a \in C \\ \emptyset & \text{otherwise.} \end{cases}$$

Notice that normative systems are not anonymous. It would certainly also be possible to consider anonymous norms, i.e. norms that are invariant under agents' names and simply forbid actions at states. However, since our main result is that non-anonymous norms are tractable on anonymous structures, we will not pay much attention to this special case in this paper. We show however, that for the example from Section 3.1, moving from anonymous to non-anonymous norms gives us increased expressive power.

3 Tractable norms for anonymous game structures

The semantic structures we will use in this section are defined as follows, following [15].

Definition 3. *A normative* 1RCGS *is a pair* $H = \langle R, \eta \rangle$ *where:*

- R *is a* 1RCGS*, and*
- η *is a normative system for* R*.*

In addition to the notions introduced for 1RCGSs, we also need access to the partial profiles which the agents in some coalition B can choose, assuming that agents in some other coalition A comply to η. The straightforward way of defining such profiles is to go via an explicit representation of compliant action-tuples for B, defined as follows.

Definition 4. *Given a* 1RCGS H*, a state* q *in* H *and two coalitions* $A, B \subseteq \mathcal{A}$*, an* $\eta \restriction A$*-compatible* B*-action at* q *is a vector* $\rho : B \to \mathbb{N}^+$ *such that:*

$$\forall b \in B \ : \ \rho(b) \in [\mathbb{A}_q] \setminus (\eta \restriction A)(q, b).$$

We let $\mathsf{Act}^{\eta}_A(q, B)$ *denote the set of all* $\eta \restriction A$*-compatible* B*-actions at* q*.*

Then B-actions give rise to B-profiles as follows.

Definition 5. *If* ρ *is an* $\eta \restriction A$*-compatible* B*-action, then the corresponding* $\eta \restriction A$*-compatible* B*-profile is a vector* $s_B^{\eta\restriction A} : [\mathbb{A}_q] \to \mathbb{N}$ *such that:*

$$\forall i \in [\mathbb{A}_q] : \left(s_B^{\eta\restriction A}\right)(i) = |\{b \in B \mid \rho(b) = i\}| \,.$$

We gather all $\eta \restriction A$*-compatible* B*-profiles for which there is a corresponding* $\eta \restriction A$*-compatible* B*-action at* q *in the set* $P^{\eta}_A(q, B)$*.*

Notice that a direct computation of this set, using Definition 4, requires computing the set $\mathsf{Act}^{\eta}_A(q, B)$, which can have exponential size in the number of agents from B. This would defeat the purpose of compact representation, the aim of which is to ensure that complexity of model checking remains polynomial in the number of agents as long as the number of actions is constant. It turns out, however, that computation of $\mathsf{Act}^{\eta}_A(q, B)$ can be avoided for *arbitrary* (non-anonymous) normative systems, and that a polynomial-time procedure can be used instead. We return to this challenge in Section 3.2, after we have defined truth on normative 1RCGS models.

To do this, we need some more notation. Given $F \in P^{\eta}_A(q, B), G \in P^{\eta}_C(q, D)$, we say that $F \geq G$ if $A = C$ and for every $i \in [\mathbb{A}_q]$ we have $F_i \geq G_i$. Given two states $q, q' \in Q$, we say that q' is a *successor* of q if there is some $F \in P(q)$ such that $\delta(q, F) = q'$. A *computation* is an infinite sequence $\lambda = q_0 q_1 \ldots$ of states such that for all positions $i \geq 0$, q_{i+1} is a successor of q_i. We follow standard abbreviations, hence a q-computation denotes a computation starting at q, and $\lambda[i]$, $\lambda[0, i]$ and $\lambda[i, \infty]$ denote the i-th state, the finite prefix $q_0 q_1 \ldots q_i$ and the infinite suffix $q_i q_{i+1} \ldots$ of λ for any computation λ and its position $i \geq 0$, respectively.

Definition 6. *An $\eta \restriction A$-compatible B-strategy is a map $s_B : Q \to \bigcup_{q \in Q} P_A^\eta(q, B)$ such that:*

$$s_B(q) \in P_A^\eta(q, B) \text{ for each } q \in Q.$$

We denote the set of all such strategies by $strat_A^\eta(B)$.

Notice that if $s \in strat_A^\eta(\mathcal{A})$ for some $A \subseteq \mathcal{A}$, then if we apply $\delta(q)$ to $s(q)$ we obtain a unique new state $q' = \delta(q, s(q))$. Iterating, we get the *induced* computation $\lambda_{s,q} = q_0 q_1 \ldots$ such that $q = q_0$ and $\forall i \ \geq 0 : \delta(q_i, (s(q_i))) = q_{i+1}$. Given $s_B \in strat_A^\eta(B)$ and a state q we get an associated *set* of computations $out(s_B, q)$. This is the set of all computations that can result when at any state, B is acting in the way specified by s_B. That is,

$$out(s_B, q) \ := \ \{\lambda_{s,q} \mid s \in strat_A^\eta(\mathcal{A}) \text{ and } s_B \leq s\}. \tag{1}$$

We can now define normative satisfaction on 1RCGS's as follows.

Definition 7. *Given a normative* 1RCGS *(H, η), a state q and a coalition $A \subseteq \mathcal{A}$, truth of φ on (H, η) under A-compliance is defined inductively.*

- $H, \eta, A, q \models p$ *iff* $q \in \pi(p)$
- $H, \eta, A, q \models \neg\varphi$ *iff* $H, \eta, A, q \not\models \varphi$
- $H, \eta, A, q \models \varphi \vee \psi$ *iff* $H, \eta, A, q \models \varphi$ *or* $H, \eta, A, q \models \psi$
- $H, \eta, A, q \models \langle\langle C \rangle\rangle \bigcirc \varphi$ *iff* $\exists s_C \in strat_A^\eta(C) : \forall \lambda \in out(s_C, q) : \lambda[1] \models \varphi$
- $H, \eta, A, q \models \langle\langle C \rangle\rangle \Box \varphi$ *iff* $\exists s_C \in strat_A^\eta(C) : \forall \lambda \in out(s_C, q) :$ $\forall i \geq 0 : \lambda[i] \models \varphi$
- $H, \eta, A, q \models \langle\langle C \rangle\rangle \varphi \mathcal{U} \psi$ *iff* $\exists s_C \in strat_A^\eta(C) : \forall \lambda \in out(s_C, q) : \exists i \geq 0 :$ $(\lambda[i] \models \psi \wedge \forall j \in [i] : \lambda[j] \models \varphi)$
- $H, \eta, A, q \models \langle B \rangle \varphi$ *iff* $H, \eta, B, q \models \varphi$

Clearly, to solve the model checking problem for this logic, we need to compute sets of the form $P_A^\eta(q, B)$, and how to do this efficiently is the main obstacle preventing a quick algorithm. We address and resolve this challenge in Section 3.2, but first we consider an example.

3.1 Example

For a simple illustration of the kind of reasoning we can perform using norms on anonymous game structures let us assume we have a system set up to perform two tasks, p_1 and p_2. Let us further assume that the system contains agents $\mathcal{A} = [n]$ where, for simplicity, we assume n is a multiple of 10. Also, assume that every agent must choose to contribute to either p_1 or p_2, a choice we encode as a choice between shared actions α_{p_1} and α_{p_2}. If the task p_1 is successfully performed, p_1 becomes true in the next state, and similarly for the task p_2.

As it happens, our system is such that in order for p_1 to be successfully performed we need $80 - 90\%$ of the agents to contribute towards p_1. That is, for p_1 to become true, such a percentage of agents have to choose α_{p_1} as their action. On the other hand, in order for p_2 to be true in the next state, we need

$20-60\%$ of the agents to perform α_{p_2}. In Figure 1 we depict an 1RCGS modelling such a scenario.[5]

Notice that if *both* p_1 and p_2 are to be performed successfully, we need precisely 20% of the agents to perform α_{p_2} while the remaining 80% choose to do α_{p_1}. It follows that in order to successfully complete both tasks, we need *coordination*. In fact, as it stands, we need everyone to coordinate their actions with everyone else. In terms of ATL, since successful completion of both p_1 and p_2 results from a *unique* profile, only the *grand coalition* can ensure $p \wedge q$. That is, while we have $H, q \models \langle\langle \mathcal{A} \rangle\rangle \bigcirc (p \wedge q)$, we also have $H, q \models [\![A]\!] \bigcirc (\neg p \vee \neg q)$ for all $A \subset \mathcal{A}$.[6]

Moreover, notice that even if some coalition $A \subseteq \mathcal{A}$ can observe what the agents in $\mathcal{A} \setminus A$ do, they might not necessarily respond in such a way that $p \wedge q$ becomes true. To see this, assume that A contains 60% of the agents. Then if the remaining agents all perform α_{p_2}, it becomes impossible for A to respond in such a way that p_1 becomes true. We have, in particular, $\langle\langle \mathcal{A} \setminus A \rangle\rangle \bigcirc \neg p$.

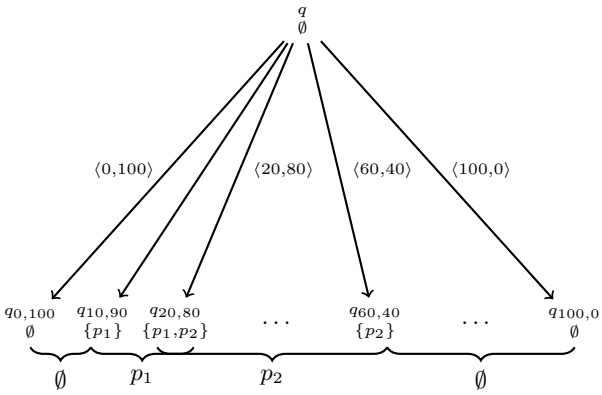

Fig. 1. A coordination problem resolved by norms

Suppose that we want to use norms to achieve $p \wedge q$ even under the assumption that only those agents that are in A are capable of coordinating their actions. Clearly, this is possible. For instance, if we simply demand that $\mathcal{A} \setminus A$ all perform the same action, and they comply, then, assuming the norm to be common knowledge, A can adapt accordingly. Somewhat more subtly, notice that in order to ensure $[\![\mathcal{A} \setminus A]\!] \bigcirc (p \wedge q)$ we do not require such a powerful norm. It is sufficient,

[5] The pairs used to decorate transitions denote profiles, with the first coordinate being the percentage of agents doing α_{p_1}, and the second coordinate being those who do α_{p_2}. We have $\delta(q_0, \langle i, j \rangle) = q_{i,j}$ for all such tuples. We omit reflexive loops for all states $q_{i,j}$.

[6] $[\![]\!]$ is the dual of the strategic ability operator $\langle\langle\rangle\rangle$. Intuitively, $[\![A]\!]\varphi$ means that coalition A can not avoid φ.

in particular, to fix some $B \subseteq \mathcal{A} \setminus A$ containing 20% of the agents, and introduce the norm η defined by:

$$\eta(q,a) = \begin{cases} \{\alpha_{p_2}\} \text{ if } q = q_0, a \in B \\ \emptyset \text{ otherwise.} \end{cases}$$

As long as B complies, A can indeed achieve $p \wedge q$ *as long as they observe* what the other agents do and adapt accordingly. In logical terms, we have $H, \eta, \emptyset, q_0 \models \langle B \rangle [\![\mathcal{A} \setminus A]\!] \bigcirc (p \wedge q)$. If it is not obvious, we leave it to the reader to verify this, possibly by using `mcheck` from Algorithm 1.

The toy example considered here also serves to illustrate that non-anonymous norms give increased expressive power compared to norms that just forbid a set of actions. Consider, in particular, the situation when we want to empower A to *choose* whether p_1 or p_2 is to become true, irrespectively of what the remaining agents do. Using a non-anonymous norm, this can be achieved by choosing $B' \in \mathcal{A}$ containing 10% of the agents such that $B' \cap A = B' \cap B = \emptyset$. To see this, consider the norm η' defined by:

$$\eta'(q,a) = \begin{cases} \{\alpha_{p_2}\} \text{ if } q = q_0, a \in B \\ \{\alpha_{p_1}\} \text{ if } q = q_0, a \in B' \\ \emptyset \text{ otherwise.} \end{cases}$$

Then, as long as $B \cup B'$ comply, we have at least 10% doing α_{p_2} and 20% doing α_{p_1}, from which it follows that p_1 is ensured as long as all members of A perform α_{p_1}, while p_2 is ensured, for instance, if 50% of the members in A perform α_{p_2}. We have, in particular, $H, \eta', \emptyset, q \models \langle B \cup B' \rangle (\langle\!\langle A \rangle\!\rangle \bigcirc p \wedge \langle\!\langle A \rangle\!\rangle \bigcirc q)$. It is not hard to see that no anonymous norm can achieve this, as long as only 30% of the agents are assumed to comply with it.

3.2 Characterizing $\eta \restriction A$-compatible B-profiles

In this section, we will provide a characterization showing that quick computation of the sets $P_A^\eta(q,B)$ is indeed possible. Towards this result, we first observe the following simple fact, the proof of which is trivial and omitted.

Whenever we use the "+" symbol with respect to vectors, we mean addition coordinate-wise.

Proposition 1. *Given a* 1RCGS*, a normative system η, coalitions $A, B \subseteq \mathcal{A}$ and a state $q \in Q$, we have $F \in P_A^\eta(q,B)$ if, and only if,*

$$\exists F_1 \in P(q, B \setminus A),\ \exists F_2 \in P_A^\eta(q, A \cap B) \text{ s.t. } F = F_1 + F_2.$$

We will also need the following auxiliary function.

Definition 8. *Given an* 1RCGS *H, a normative system η and any state $q \in Q$ we define, for all $E \subseteq \mathbb{A}_q, A \subseteq \mathcal{A}$, the following set:*

$$\mathcal{C}_\eta^q(E,A) = |\{x \in A \mid \eta(q,x) \cap E \neq \emptyset\}|.$$

So $\mathcal{C}^q_\eta(E, A)$ returns the *number* of agents in A that have a legal action in E at q. Using this function allows us to characterize $P^\eta_A(q, B)$ more compactly using a matching argument, giving rise to the following lemma, towards tractable model checking.

Lemma 1. *For any* 1RCGS*, any normative system* η *and any* $A, B \subseteq \mathcal{A}$ *we have* $F \in P^\eta_A(q, B)$ *iff* $F = F_1 + F_2$ *for some* $F_1 \in P(q, B \setminus A)$ *and some* $F_2 \in P(q, A \cap B)$ *such that:*

$$\forall E \subseteq [\mathbb{A}_q] : \mathcal{C}^q_\eta(E, A \cap B) \geq \sum_{i \in E} F_2(i). \tag{2}$$

Proof. $\Rightarrow$) Trivial.
$\Leftarrow$) Assume that we have $F = F_1 + F_2$ for $F_1 \in P(q, B \setminus A)$ and $F_2 \in P(q, A \cap B)$ such that (2) holds. We demonstrate existence of $\rho \in \mathsf{Act}^\eta_A(q, B)$ that induces the profile F_2, i.e., such that

$$F_2(i) = |\{x \in A \cap B \mid \rho(x) = i\}| \text{ for all } i \in [\mathbb{A}_q].$$

We will think of ρ as the solution of a matching problem in a bipartite graph: Let $G = (V_1, V_2, E)$ where $V_1 = A \cap B$, $V_2 = \{i_j \mid i \in [\mathbb{A}_q], j \in [F_2(i)]\}$ are the two sets of nodes and $E = \{(x, i_j) \mid i \notin \eta(q, x)\}$ is the set of edges. Notice that $|V_1| = |V_2|$ since $F_2 \in P(q, A \cap B)$, and that the graph is indeed bipartite. For all subsets of $V \subseteq V_2$, let $V^- = \{x \in V_1 \mid \exists i_j \in V_2 : (x, i_j) \in E\}$. Then, since F_2 satisfies (2), it follows that for all $V \subseteq V_2$ we have $|V^-| \geq V$. This means that the conditions of Hall's marriage theorem are all fulfilled (well known from graph theory, originally published in [9]), meaning that there exists a set $E' \subseteq E$ such that for every $i_j \in V_2$ there is a unique $x \in V_1$ such that $(x, i_j) \in E'$, i.e., such that E' is a matching in G. Let us define the vector $\rho : A \cap B \to \mathbb{N}^+$ such that $\rho(x) = i$ for all $(x, i_j) \in E'$. Clearly, since E' is a matching, this is well-defined and we have $\rho \in \mathsf{Act}^\eta_A(q, B)$ as desired. Moreover, it is easy to see that ρ corresponds to F_2 in the sense of Definition 5. We conclude that $F_2 \in P^\eta_A(q, B \cap A)$. Then, from Proposition 1 it follows that $F = F_1 + F_2 \in P^\eta_A(q, B)$, concluding the proof. □

In Figure 2 we illustrate how $P^\eta_A(q, B)$ is generated, by calculating the sets $P(q, B \setminus A)$ and $P(q, B \cap A)$ the latter of which is then restricted to the elements which satisfy Condition (2) (in Lemma 1). The parameters of the situation illustrated is the number of actions $\mathbb{A}_q = 3$ and the set of agents $\mathcal{A} = \{a, b, c, d, e\}$ with agents $A = \{b, c, d\}$ complying to η, and the agents for which we are making a set of profiles for are contained in $B = \{c, d, e\}$. The consequence of η for the two agents in B which do comply, is that it forbids action 2 for agent c and actions 1 and 2 for agent d.

In light of Lemma 1, it is clearly possible, as long as the number of actions is constant, to generate $P^\eta_A(q, B)$ in polynomial time for all A, B, q. We simply run through all $F \in P(q, A \cap B)$ and check if Condition (2) holds. This involves running though all subsets of $\mathbb{A}_q$, but still it only requires a constant number of traversals of $A \cap B$. Then the set $P^\eta_A(q, B)$ is obtained from any such F passing

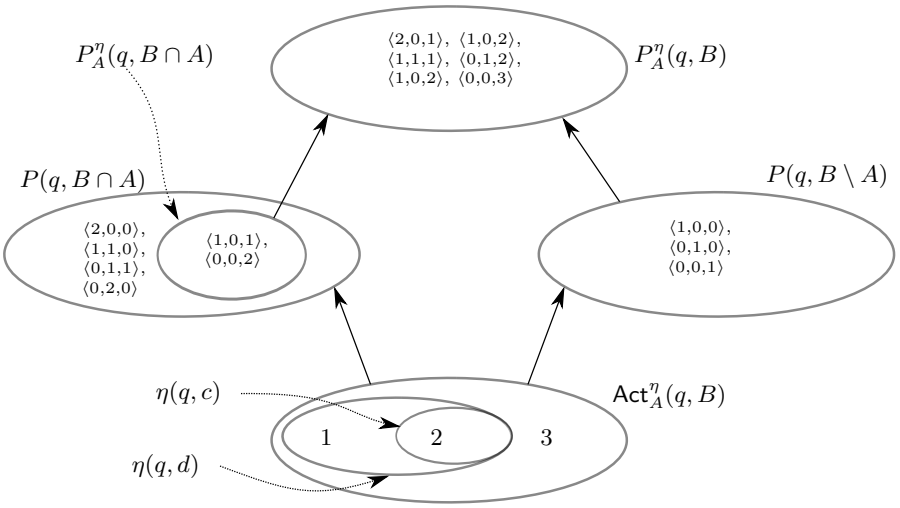

Fig. 2. Illustration of Lemma 1.

the test, when added to any vector from the set $P(q, B \setminus A)$, as detailed in Algorithm 3.

We mention that the construction in the proof of Lemma 1 mirrors the construction used in [6] to establish that finding pure Nash equilibria in an anonymous normal form game is decidable in polynomial time provided the number of actions remain constant. This result, in particular, is also obtained by an application of Hall's marriage theorem.

More importantly, given an 1RCGS H, a normative system η, a state $q \in Q$ and coalitions A, B, it seems clear that we can define an anonymous normal form game such that $P_A^\eta(q, B)$ is the set of pure Nash equilibria in this game. We omit the details due to space restrictions, but remark that as Lemma 1 can be seen as a corollary of results from [6], it follows that computing $P_A^\eta(q, B)$ can also be done by employing the more subtle techniques introduced there, used to prove membership in the complexity class TC^0. This means, in particular, that the algorithm presented in the next section, while showing that model checking is tractable, could be improved on this point. Here, however, we do not focus on the design of optimal procedures, but on clearly conveying the main result and the ideas that have precipitated it.

3.3 Tractable model checking

The algorithm for checking truth of φ in a normative 1RCGS follows exactly the same pattern as the standard model checking algorithm used to do model checking on CGS models, see e.g., [11]. Given a CGS model S and a formula φ,

this algorithm processes φ recursively and returns the set of states $q \in Q$ where φ is true. To deal correctly with $\langle\langle A\rangle\rangle\Box\varphi$ and $\langle\langle A\rangle\rangle\varphi\mathcal{U}\psi$ the algorithm relies on the following fixed point characterizations, which are well-known to hold for ATL, see for instance [11], and are also easily seen to be true on any normative 1RCGS model, c.f., Definition 7:

$$\begin{aligned} \langle\langle A\rangle\rangle\Box\varphi &\leftrightarrow \varphi \wedge \langle\langle A\rangle\rangle \bigcirc \langle\langle A\rangle\rangle\Box\varphi \\ \langle\langle A\rangle\rangle\varphi_1\mathcal{U}\varphi_2 &\leftrightarrow \varphi_2 \vee (\varphi_1 \wedge \langle\langle A\rangle\rangle \bigcirc \langle\langle A\rangle\rangle\varphi_1\mathcal{U}\varphi_2 \end{aligned} \tag{3}$$

In light of this, the correctness of the algorithm `mcheck`, shown in Algorithm 1, follows trivially if we can establish correctness of the algorithm `enforce`, shown in Algorithm 2.

Algorithm 1 `mcheck`(H, η, A, φ) algorithm

if $\varphi = p \in \Pi$ **then**
  return $\pi(p)$
if $\varphi = \neg\psi$ **then**
  return $Q \setminus$ `mcheck`(H, η, A, ψ)
if $\varphi = \psi \vee \psi'$ **then**
  return `mcheck`$(H, \eta, A, \psi) \cup$ `mcheck`(H, η, A, ψ')
if $\varphi = \langle\langle B\rangle\rangle \bigcirc \psi$ **then**
  return $\{q \mid$ `enforce`$(H, \eta, A, q, B,$ `mcheck`$(H, \eta, A, \psi))\}$
if $\varphi = \langle\langle B\rangle\rangle\Box\psi$ **then**
  $Q_1 := Q$
  $Q_2 :=$ `mcheck`(H, η, A, ψ)
  while $Q_1 \not\subseteq Q_2$ **do**
    $Q_1 := Q_2$
    $Q_2 := \{q \in Q \mid$ `enforce`$(H, \eta, A, B, q, Q_2)\} \cap Q_2$
  return Q_1
if $\varphi = \langle\langle B\rangle\rangle\psi\mathcal{U}\psi'$ **then**
  $Q_1 := \emptyset$
  $Q_2 = mcheck(H, \psi)$
  $Q_3 = mcheck(H, \psi')$
  while $Q_3 \not\subseteq Q_1$ **do**
    $Q_1 := Q_1 \cup Q_3$
    $Q_3 := \{q \in Q \mid$ `enforce`$(H, \eta, A, B, q, Q_1)\} \cap Q_2$
  return Q_3
if $\varphi = \langle A'\rangle\psi$ **then**
  return `mcheck`(H, η, A', ψ)

This algorithm answers, given a normative 1RCGS (H, η), a state q, coalitions $A, B \subseteq \mathcal{A}$ and a set of states Q', whether or not there is some strategy $s_B \in strat_A^\eta(B)$ such that $\{\lambda[1] \in Q \mid \lambda \in out(s_B, q)\} \subseteq Q'$. Clearly, such a strategy exists if, and only if, there is some $F_B \in P_A^\eta(q, B)$ such that for all $F \in P_A^\eta(q, \mathcal{A})$, if $F_A \leq F$ then $\delta(q, F) \in Q'$. Thus, correctness of `enforce` follows if the algorithm `comp`, shown in Algorithm 3, correctly computes the necessary

Algorithm 2 `enforce`(H, η, Aq, B, Q') algorithm

$S_{pro} = \texttt{comp}(\eta, A, q, B)$ // $S_{pro} = P_A^\eta(q, B)$
$S_{ant} = \texttt{comp}(\eta, A, q, B)$ // $S_{ant} = P_A^\eta(q, (\mathcal{A} \setminus B))$
for $F_B \in S_{pro}$ **do**
  $x = true$
  for $F_{B'} \in S_{ant}$ **do**
    if $\delta_q(F_B + F_{B'}) \notin Q'$ **then**
      $x = false$
  if $x = true$ **then**
    return true
return false

sets $P_A^\eta(q, B)$. This, in turn, clearly follows from Lemma 1. To see this, notice that the step when we place agents in the set T corresponds exactly to the calculation of $\mathcal{C}_\eta^q(E, A)$.[7]

Moreover, notice that all of the procedures involved in model checking have polynomial complexity in the length of the formula and the size of the model. This follows by the fact that the sizes of S_{pro} and S_{ant}, used by `enforce` and calculated by `comp`, have sizes bounded above by $\frac{(|B|+(|\mathbb{A}_q|-1))!}{|B|!(|\mathbb{A}_q|-1)!}$ and $\frac{(|\mathcal{A}\setminus B|+(|\mathbb{A}_q|-1))!}{|\mathcal{A}\setminus B|!(|\mathbb{A}_q|-1)!}$ respectively. These combinatorial expressions are both bounded above by $|\mathcal{A}|^{|\mathbb{A}_q|}$, so there is indeed no exponential dependence on the number of agents, only on the number of actions. Also remember that we compute S_{pro} and S_{ant} effectively, by applying Lemma 1. The main result follows.

Theorem 1. *Given a normative* 1RCGS (H, η)*, a state* $q \in Q$*, a coalition* $A \subseteq \mathcal{A}$ *and a formula* φ*: Deciding if* $H, \eta, A, q \models \varphi$ *takes polynomial time in the size of* H *and the length of* φ*.*

4 Conclusion

In this paper, we have considered concurrent game structures that satisfy anonymity. Following [15,14], we represent these structures compactly, avoiding models that have exponential size in the number of agents. Then we consider normative systems applied to such models, resulting in the logic NCHATL. Our main technical result is that this logic still admits a tractable algorithm for the model checking problem.

More generally, we believe our work serves to establish interesting connections, both conceptual and technical, between recent work in algorithmic game theory and recent work on logics for strategic ability of coalitions of agents. It seems, in particular, that a major challenge which is becoming increasingly important to both these fields is the need for compact representations, allowing us

[7] This implementation of the $P_A^\eta(q, B)$, collecting agents in T, could be optimized if we just count the first occurrence of a satisfying condition (where a is added to T) and move on to the next agent. We use a set to simplify the presentation.

Algorithm 3 $\texttt{comp}(\eta, A, q, B)$ algorithm including $T = \mathcal{C}^q_\eta(E, A)$

```
R = P(q, B \ A)
for x ∈ P(q, B ∩ A) do
  y = true
  for E ⊆ 𝔸_q do
    T = ∅
    for a ∈ B ∩ A do
      for α ∈ E do
        if α ∉ η(q, a) then
          add a to T
  if |T| < Σ_{α∈E} x_α then
    y = false
  if y = true then
    add x to R
return R
```

to make use of established formalisms to analyse systems with a large number of participating agents.

In order for this to become feasible in practice, we certainly require representations and notions that avoid introducing exponential time-dependence on the number of agents that are present. The danger, however, is that when formulating restrictions that make this possible, one deprives the underlying formalism of crucial expressive power. In this paper, we have addressed this worry for ATL, and shown that norms can be used to regain some of what is lost by requiring anonymity.

Moreover, and somewhat surprisingly, it turns out that even non-homogeneous norms can be implemented without introducing any exponential dependence on the agents. We conclude, therefore, that normative systems are a good candidate in general for giving compact multi-agent formalism a limited, but useful, means for talking about such heterogeneous properties that can be expressed without resulting in an exponential blow-up of crucial decision problems.

References

1. T. Ågotnes, W. van Der Hoek, J. A. Rodríguez-Aguilar, C. Sierra, and M. Wooldridge. On the Logic of Normative Systems. In *Proc. of the 20th Int. Joint Conf. on Artificial Intelligence (IJCAI 07)*, pages 1175–1180, 2007.
2. T. Ågotnes, W. van der Hoek, and M. Wooldridge. Robust normative systems and a logic of norm compliance. *Logic Journal of the IGPL*, 18(1):4–30, 2009.
3. R. Alur, T. Henzinger, F. Mang, S. Qadeer, S. Rajamani, and S. Tasiran. Mocha: Modularity in model checking. In A. Hu and M. Vardi, editors, *Computer Aided Verification*, volume 1427 of *Lecture Notes in Computer Science*, pages 521–525. Springer Berlin Heidelberg, 1998.
4. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM (JACM)*, 49(5):672–713, 2002.

5. M. Blonski. Characterization of pure strategy equilibria in finite anonymous games. *Journal of Mathematical Economics*, 34:225–233, 2000.
6. F. Brandt, F. Fischer, and M. Holzer. Symmetries and the complexity of pure nash equilibrium. *Journal of Computer and System Sciences*, 75(3):163–177, 2009.
7. C. Daskalakis and C. Papadimitriou. Computing equilibria in anonymous games. In *Foundations of Computer Science, 2007. FOCS '07. 48th Annual IEEE Symposium on*, pages 83–93, 2007.
8. P. Dellunde. On the multimodal logic of normative systems. In J. S. Sichman, J. A. Padget, S. Ossowski, and P. Noriega, editors, *COIN*, volume 4870 of *Lecture Notes in Computer Science*, pages 261–274. Springer, 2007.
9. P. Hall. On representatives of subsets. *Journal of the London Mathematical Society*, s1-10(1):26–30, 1935.
10. W. Hoek, M. Roberts, and M. Wooldridge. Social laws in alternating time: effectiveness, feasibility, and synthesis. *Synthese*, 156(1):1–19, 2007.
11. W. Jamroga. Easy yet hard: Model checking strategies of agents. In M. Fisher, F. Sadri, and M. Thielscher, editors, *Computational Logic in Multi-Agent Systems*, pages 1–12. Springer-Verlag, Berlin, Heidelberg, 2009.
12. C. H. Papadimitriou. The Complexsity of Finding Nash Equilibria. In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, pages 29–52. Cambridge University Press, 2007.
13. M. Pauly. A Modal Logic for Coalitional Power in Games. *Journal of Logic and Computation*, 12(1):149–166, Feb. 2002.
14. T. Pedersen and S. Dyrkolbotn. Agents homogeneous: A procedurally anonymous semantics characterizing the homogeneous fragment of ATL. To appear in the *Proceedings of PRIMA 2013*, LNAI, 2013.
15. T. Pedersen, S. Dyrkolbotn, P. Kaźmierczak, and E. Parmann. Concurrent game structures with roles. In F. Mogavero, A. Murano, and M. Y. Vardi, editors, Proceedings 1st International Workshop on *Strategic Reasoning,* Rome, Italy, March 16-17, 2013, volume 112 of *Electronic Proceedings in Theoretical Computer Science*, pages 61–69. Open Publishing Association, 2013.
16. Y. Shoham and M. Tennenholtz. On the synthesis of useful social laws for artificial agent societies. In *Proceedings of the tenth national conference on Artificial intelligence*, AAAI'92, pages 276–281. AAAI Press, 1992.
17. Y. Shoham and M. Tennenholtz. On social laws for artificial agent societies: Offline design. *Artificial Intelligence*, 73:231–252, 1995.